LookDeep

2025-06-26

A Deep Look at Continuous Patient Monitoring

Paolo Gabriel, PhD

Be Present for Every Patient at Every Moment

About me

[PHL, JPN, USA-(OH,CA)]

Paolo Gabriel

paologabriel.com

UC San Diego, 2013 - 2019
- Ph.D. ECE - Medical Devices & Systems
- Translational Neuroengineering Lab

LookDeep Health, 2019 -
- Sr. Computer Vision Engineer
- **AI-assisted patient monitoring**

LookDeep

Recording units

Continuous patient monitoring with AI:
real-time analysis of video
in hospital care settings
(Front. Imaging, 09 March 2025)

Agenda

1. Continuous patient monitoring
2. CV development for hospital settings
3. Lessons learned
4. Discussion

Disclaimer:
This is a talk about **simple and transparent tools**, on top of a **robust system**

The need for patient monitoring

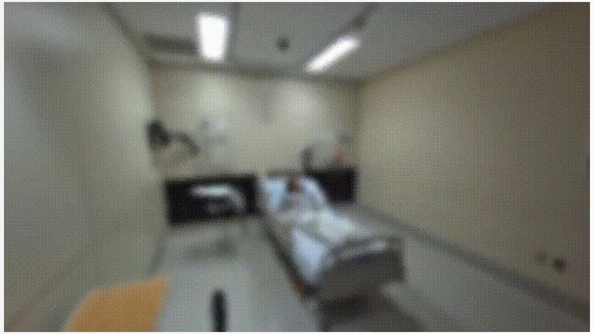

If you are staying in a hospital…

- Most of the time, you are alone and unattended
- Your status is checked on a schedule
- *Calling for assistance takes effort* ~out of scope~

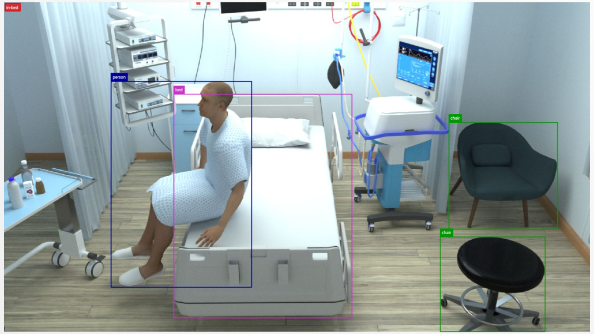

Even **simple information** can be useful

- Where is the patient, what state are they in?
- Is there staff in the room?
- *What's the diagnosis?* ~out of scope~

Before applying novel ML, **can you demonstrate the basics?**

Patient monitoring with computer vision

Why use **computer vision**?

- Direct observation is **limited**, annotation is **time consuming**
- Analyze video over extended periods with **computer vision**
  - Existing work ~ (Chen et al., 2018), (Wang et al., 2018), (Peterson et al., 2021)
- **Baseline** architecture and performance: (Gabriel et al., 2025)
  - RGB @ **1fps** on rknn NPU
  - **Yolo v4** object detection + **Farneback dense optical flow**
  - Almost **3 years** of recording at **11 hospitals**

Benchmarks for AI-driven patient monitoring, data-driven insights into **patient behavior** and **interactions**.

Recording unit

Real hospital settings

Goal: **directly observe** patients in the **noisy** clinical environment.

You want **clean setup…**

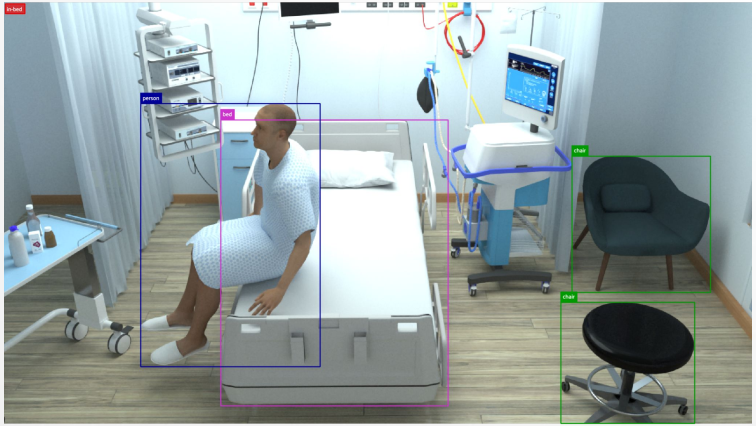

You get something **“wild” instead**

> We built and **validated** a **computer vision <u>platform</u>** for **real hospital settings!**

Examples of our **domain** (blurred for privacy)

Camera placement varies

The **LookDeep** Virtual Care Platform

How we monitor each patient **24/7**

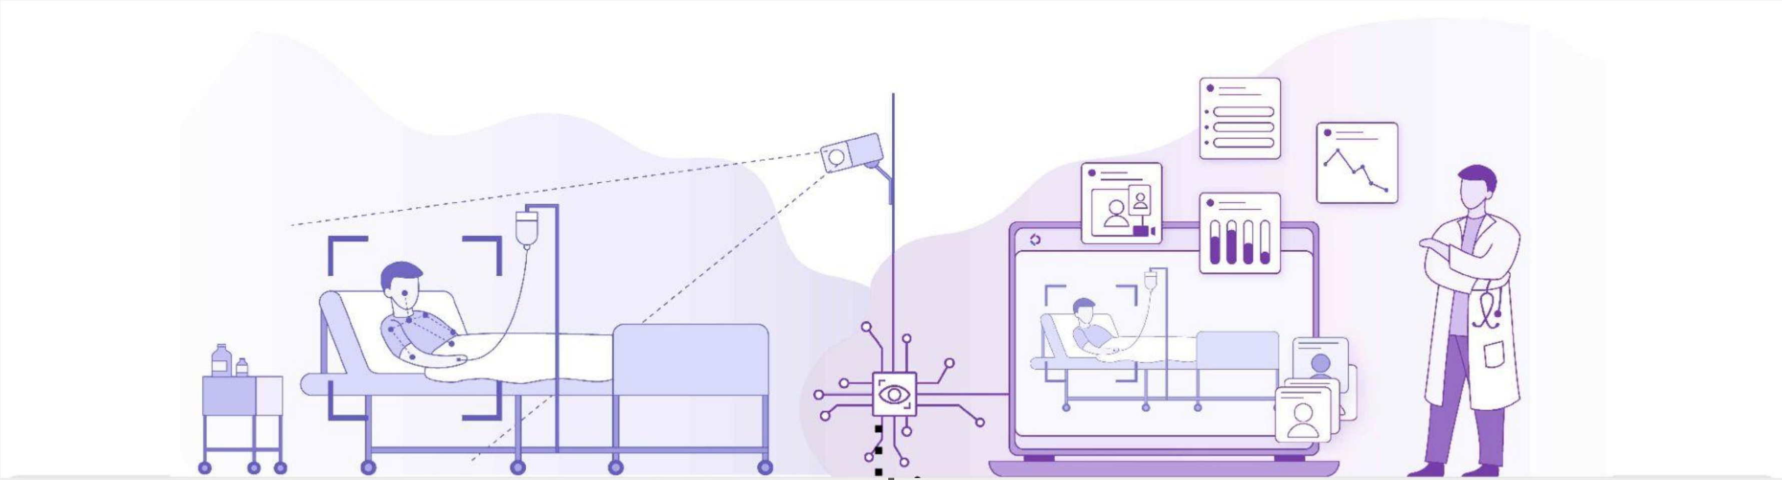

Computer Vision Pipeline

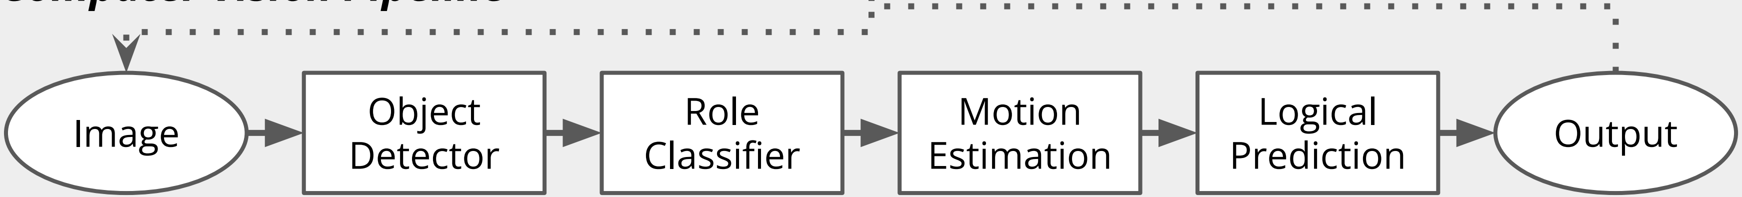

Example **output data**

Real-time pipeline

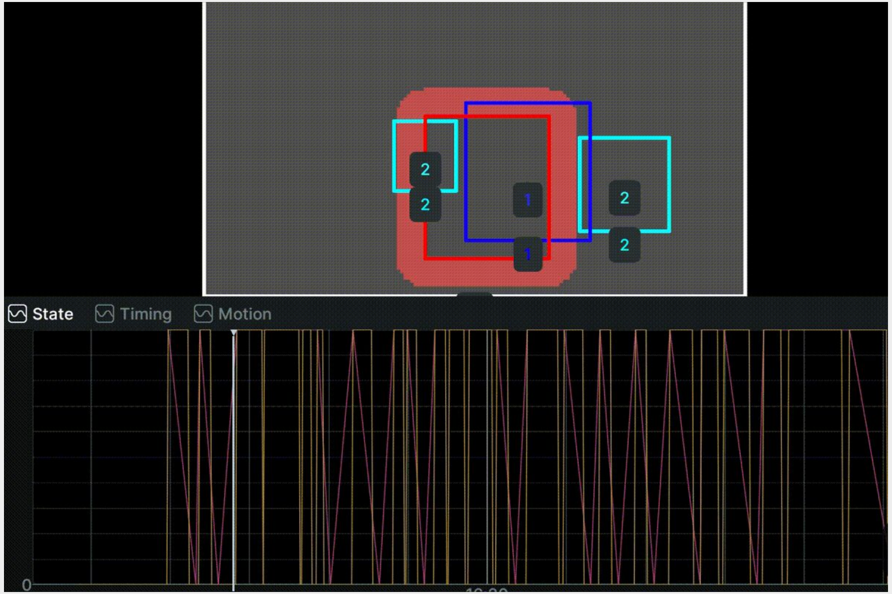

Objects, masks, motion, state changes

Time granularities (s, m, h)

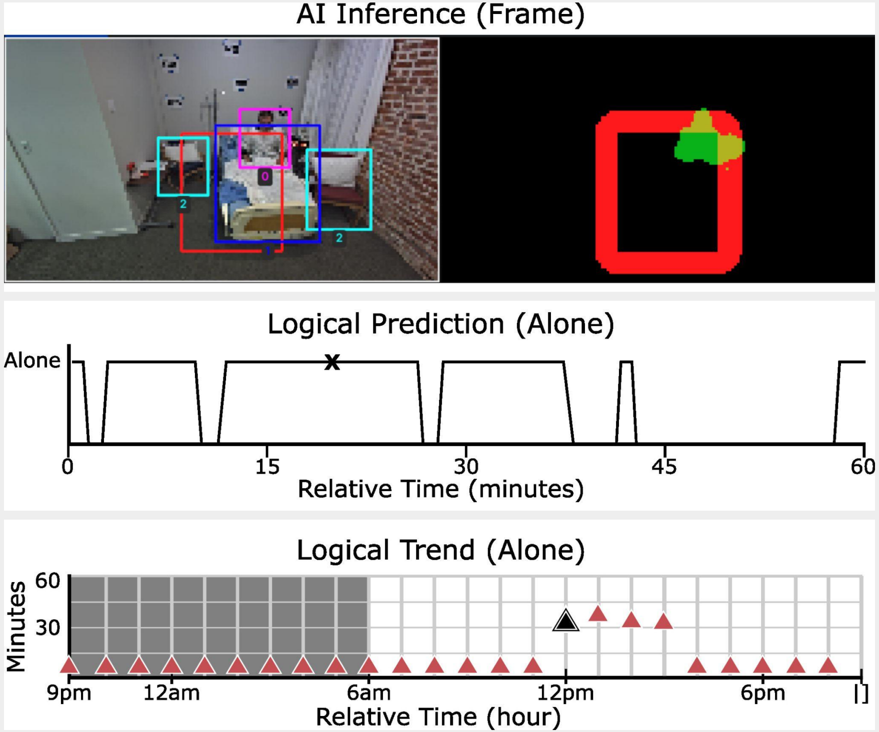

Example of our **frame-level labels**

40K+ frames at time of publication

Labeled image

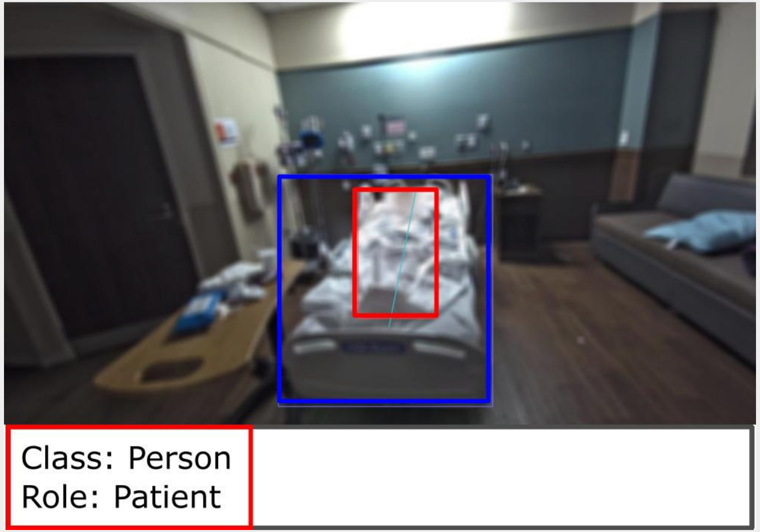

Data lake (one of several)

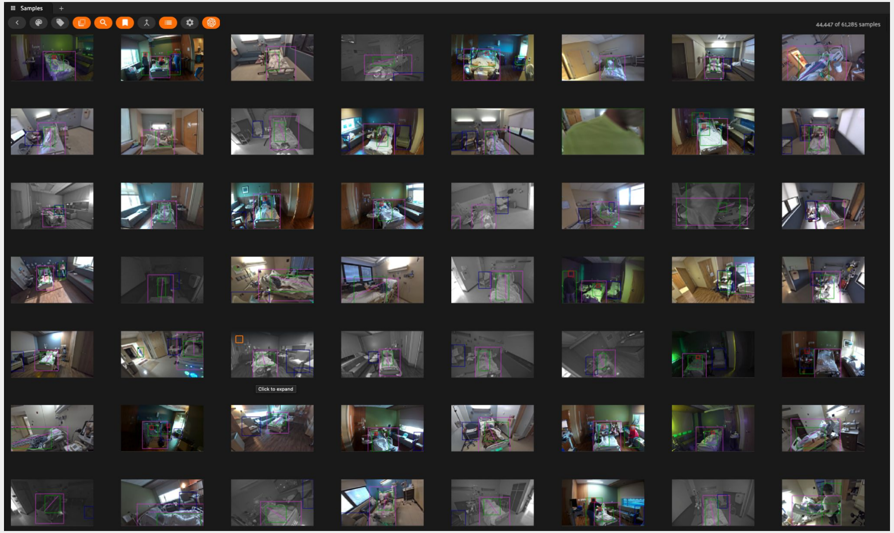

*Labeled **metadata** (e.g. "bad image", "truncated") used to **curate training** data*

Data **labeling** for **ongoing performance** management

How we **use our labels**

30K hours per month (100K+ / mo by 2025)

FiftyOne management

Image labeling

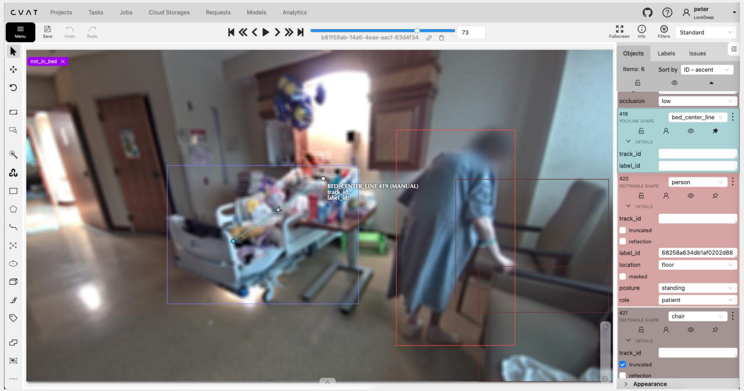

All labeling is blurred
(final image is face blurred)

Offline evaluation

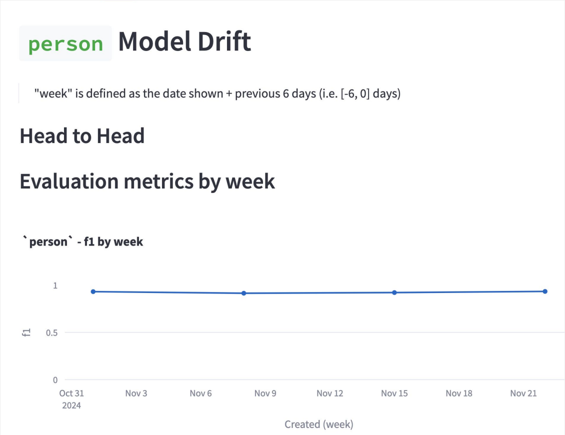

- Test set - **every 4th week**
- Compare **old vs new** models

Frame-level analysis - object detection, classification

AI Inference

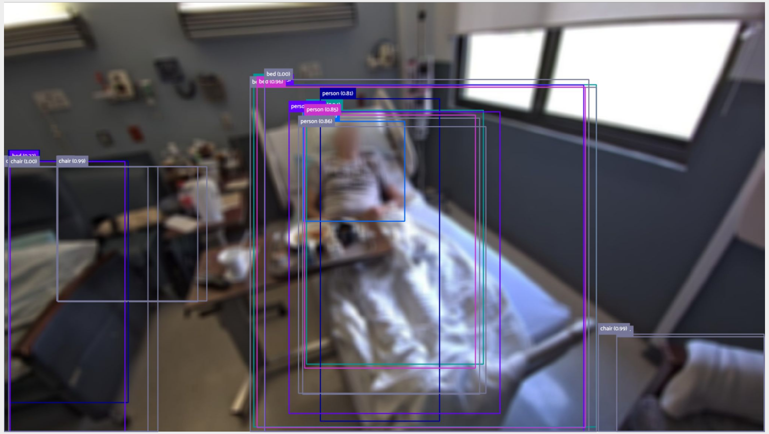

"All" objects, over time

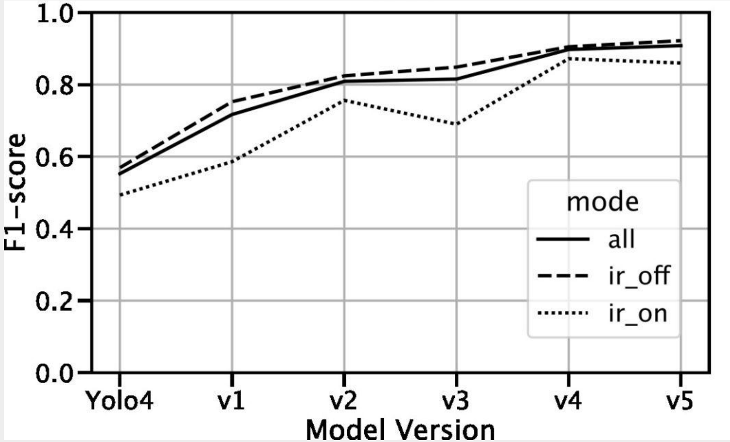

Detection and classification metrics

Model version	Object detection (person)		Role classification (patient F1)	"Patient alone" classification (F1)
	Precision	F1		
YOLOv4 (baseline)	0.98	0.41	n/a	0.28
Model v5 (2024-Q2)	0.96	0.91	0.98	0.92

Model development **over time**

Re-training models with **more data,** use **most recent test set**

"All" objects, over time

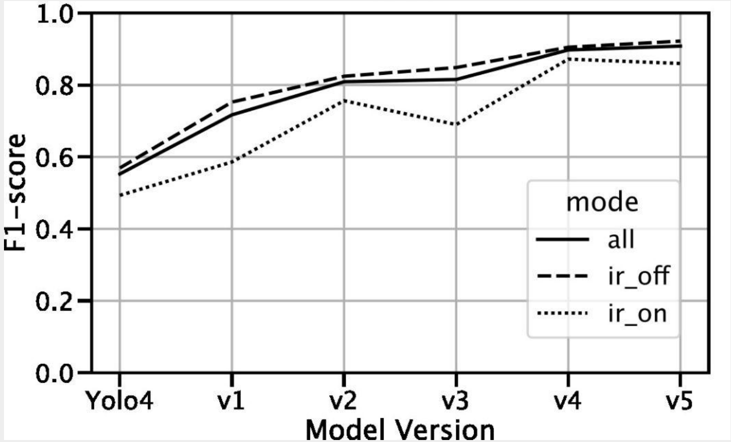

Model version	Fine-tuning data size
YOLOv4 (baseline)	n/a
Model v1 (2022 Q1)	+700
Model v2 (2023 Q2)	+2,474
Model v3 (2023 Q3)	+10,133
Model v4 (2024 Q1)	+28,914
Model v5 (2024 Q2)	+34,239

Improvement with more targeted data (e.g. patient standing at night)

Example of our **trend-level** data

Time segment labels

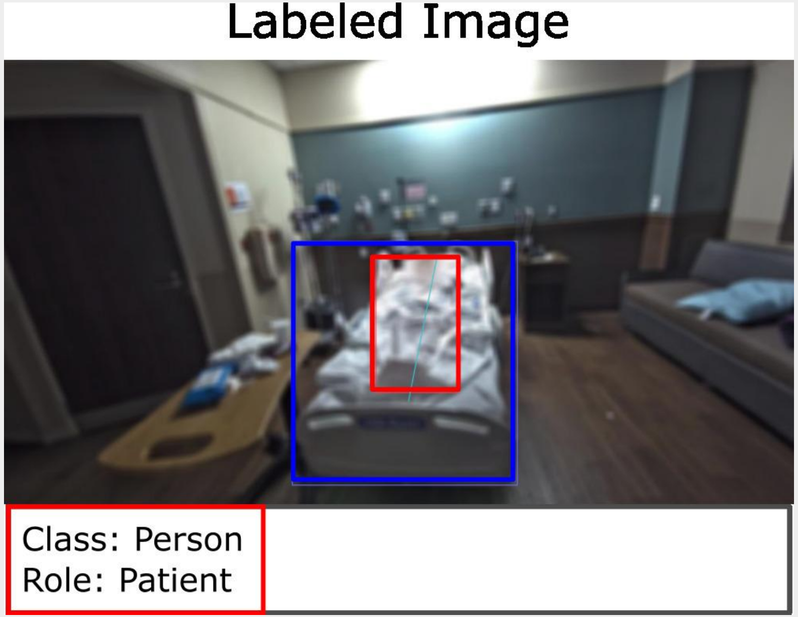

Time segment predictions

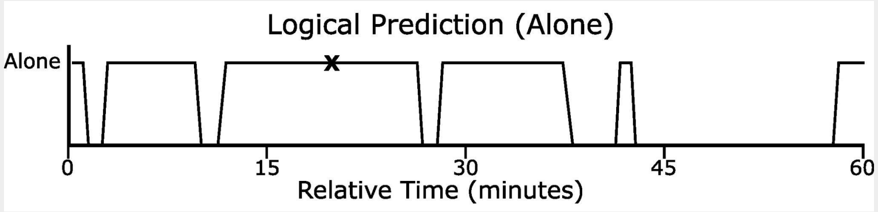

- Evaluate **consistency** of signal over time
- **Separate** logical algorithms from core CV
- Requires **video**

Trend-level analysis - “patient is alone”

Time segment predictions

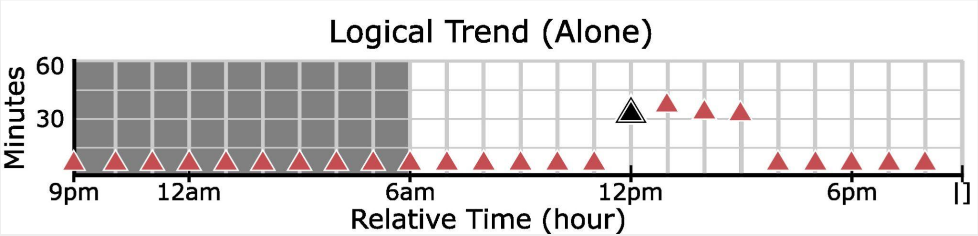

average logistic regression/manual accuracy
- **0.82 ± 0.15** across all times

Model v3, 10 patients

Compare against GT trends

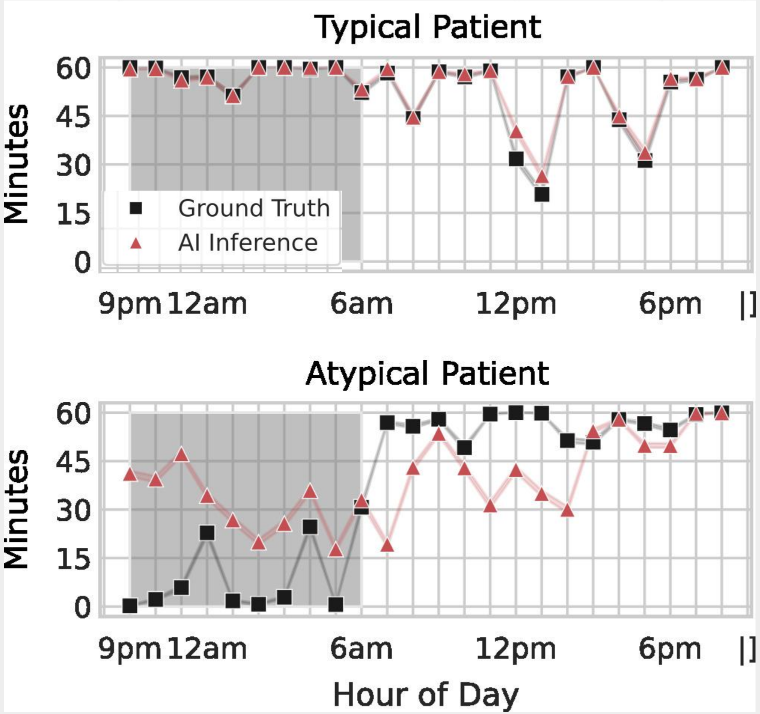

Manuscript contributions

Continuous patient monitoring with AI:
real-time analysis of video in hospital care settings
(Front. Imaging, 09 March 2025)

- AI-driven patient monitoring **system**
- **Multi-year** data collection
- **Model training** and **evaluation** process
    - Object detection, role/state classification
    - “Patient alone” trends
- Anonymized **dataset of hourly trends ->**

Public dataset

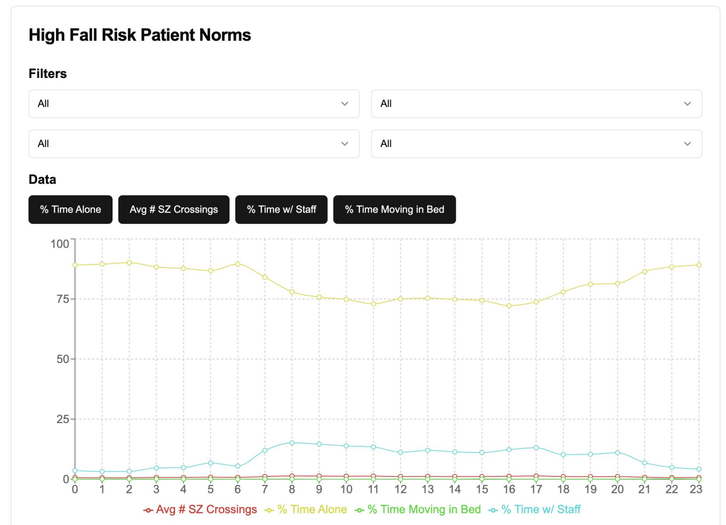

Lessons learned - **evolving data coverage**

Datasets timeline

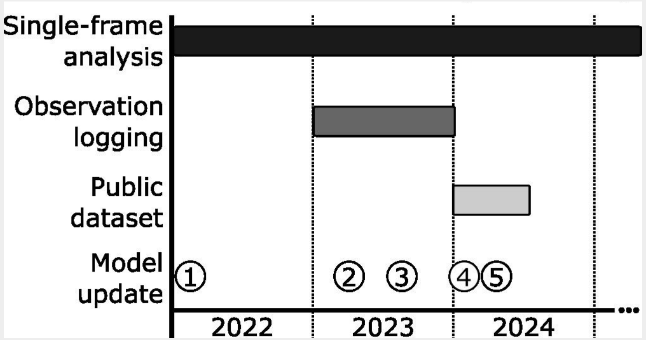

- More data over time is a **good thing**
- Have a **consistent** test set (ours is now 10k+)
- Have a **tight, continuous integration**

Demographic information

- **Metadata** enables audits
- Understand **what is there**
- Anticipate **biases**

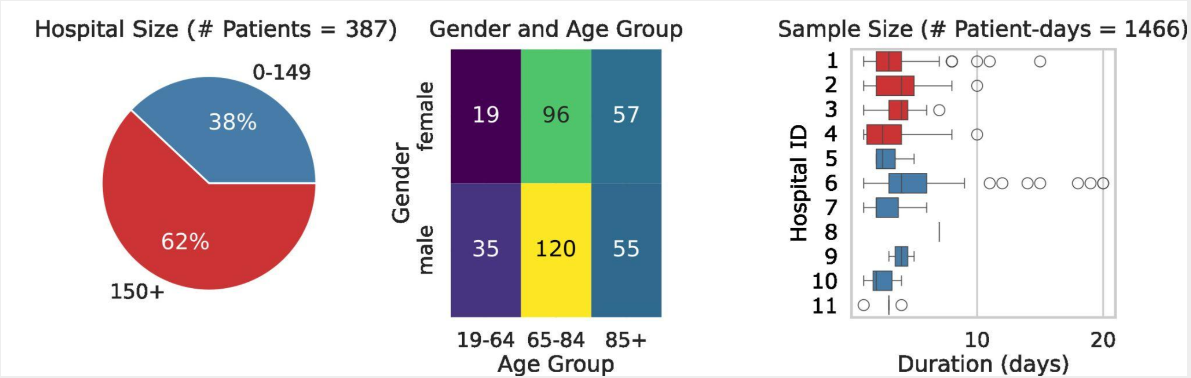

Lesson learned - generalizing across **camera conditions**

Face-blurred vs raw images

- live -> unblurred
- label -> full blur
- train -> face blur

+0.04 Δ f1-score
object detection

IR on/off

Shift data collection for extra "ON" samples

Variations in camera placement

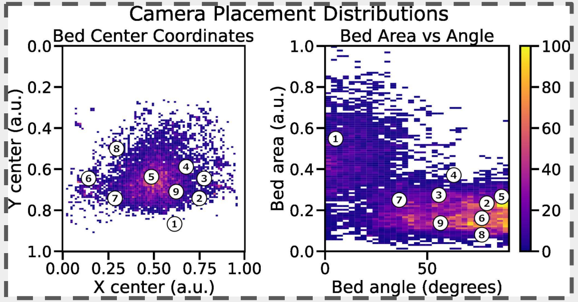

Downstream stability

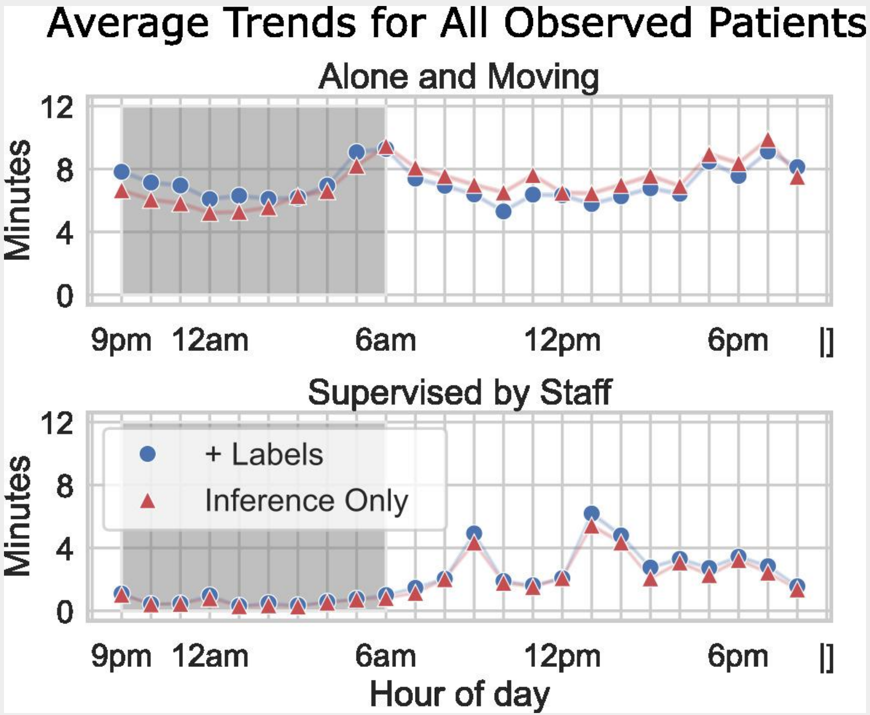

average error of 1–2 min per hour

Lessons learned - our stack of **AI data tools**

FiftyOne - image curation and analysis

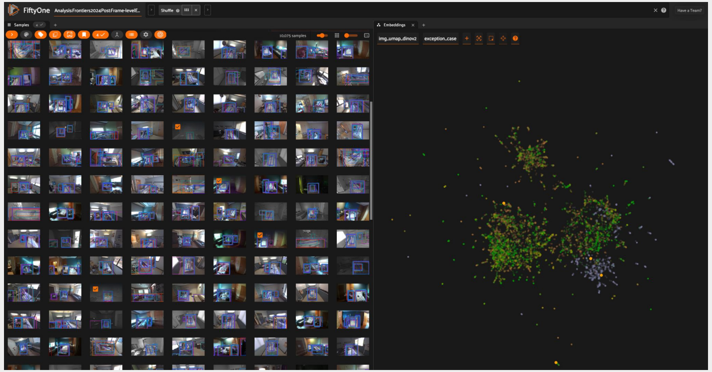

Rerun - physical data viewer

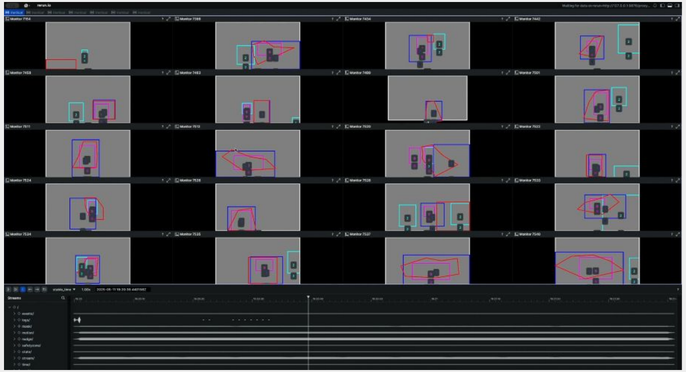

CVAT - labeling

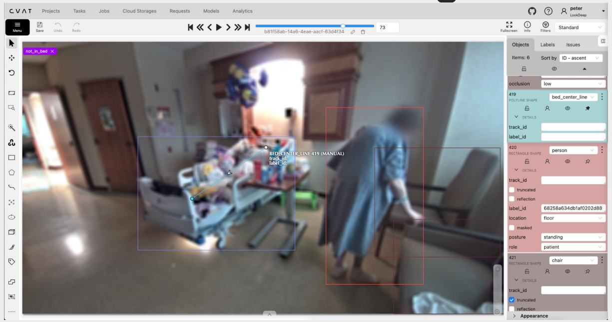

Custom tools - 3d render

GCP - data storage

Summary

- **Applied computer vision,** rigorous eval
- **11 hospitals**, over **300 high-risk fall patients**, more than **1,000 days** of inference
- **Open access** to “hour-level” features from 2024
- **Lessons learned**:
    - tag, tag, tag
    - expect imperfect conditions
    - adopt existing AI tools

We’re getting better…

Paper (v5)

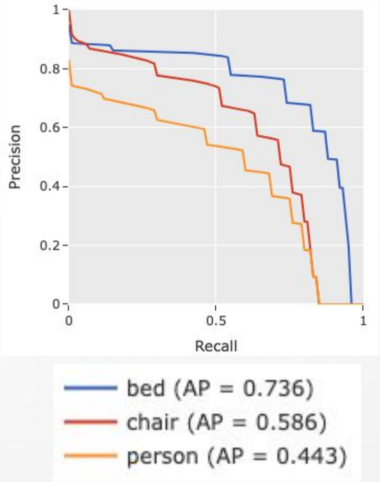

v7

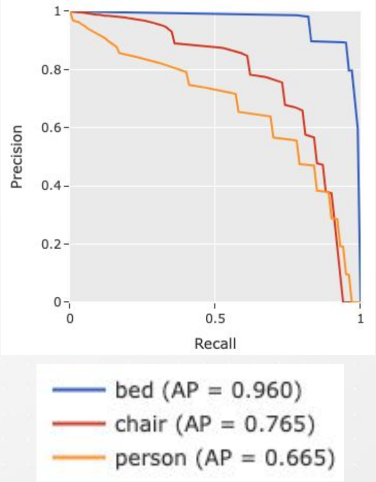

Acknowledgements

Co-authors: Peter Rehani, Tyler Troy, Tiffany Wyatt, Michael Choma, Narinder Singh

Co-workers: Guram Kajaia, James Eitzman, Bill Mers, Mike O'Brien, Jan Marti, Laura Urbisci, Tom Hata

Reviewers: Aashish P., Jacob H., Kenny C., Tejaswy P., Quirine vE., Hina S., Nazreen P.M., Sandeep K.M.

AI-agent: GPT-4-turbo for refining manuscript text, to improve clarity and organization of the research

Study participants:

> *The data used in this retrospective study was collected from patients admitted to one of eleven hospital partners across three different states in the USA. Patients provided written informed consent for monitoring as part of their standard inpatient care. To ensure patient privacy, all visual data was blurred and no identifiable information is presented.* **Thank you for participating.**

Thank you!

Questions?

paologabriel.com